

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
1 March 2001 (01.03.2001)

PCT

(10) International Publication Number
WO 01/14961 A2

(51) International Patent Classification⁷: **G06F 9/00**

(21) International Application Number: **PCT/US00/24336**

(22) International Filing Date: **28 August 2000 (28.08.2000)**

(25) Filing Language: **English**

(26) Publication Language: **English**

(30) Priority Data:
60/150,766 26 August 1999 (26.08.1999) US
60/210,334 12 June 2000 (12.06.2000) US

(71) Applicant: **PARABON COMPUTATION [US/US];**
Suite 1000, 3930 Walnut Street, Fairfax, VA 22030 (US).

(72) Inventors: **ARMENTROUT, Steven, L.**; 11879 St. Trinians Ct., Reston, VA 20191 (US). **O'CONNOR, James**; 11625 Ayreshire Road, Oakton, VA 22124 (US). **GANNON, James**; 3817 Inverness Road, Fairfax, VA 22030 (US). **SLETTEN, Brian**; Apt. 906, 900 N. Stuart Street, Arlington, VA 22203 (US). **CIER, Sean**; Apt.

12, 11800 Federalist Way, Fairfax, VA 22030 (US). **CARLSON, Sarah**; 13236 Memory Lane, Fairfax, VA 22033 (US). **DAVIS, Jonathan**; Apt. 312, 202 Columbia Road NW, Washington, DC 20009 (US). **DUPERTUIS, Greg**; 7888 Leroux Lane, Manassas, VA 20112 (US). **MCCLOUGHLIN, Scott**; 2827 28th Street, #14, Washington, DC 20001 (US). **DAVIES, Antony**; 11733 Stuart Mill Road, Oakton, VA 22124 (US).


(74) Agents: **ROBERTS, Jon, L. et al.**; Roberts Abokhair & Mardula, LLC, Suite 1000, 11800 Sunrise Valley Drive, Reston, VA 20191 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European

[Continued on next page]

(54) Title: **SYSTEM AND METHOD FOR THE ESTABLISHMENT AND UTILIZATION OF NETWORKED IDLE COMPUTATIONAL PROCESSING POWER**



(57) Abstract: A distributed computing platform using the idle computational processing power of a plurality of provider computers is disclosed. At least one networked server collects tasks from client computers, schedules and distributes the tasks to networked provider computers, and collects and returns results to client computers. A client API forms tasks and collects results. A compute engine operates on the provider computers to communicate with the server and execute tasks using idle computational power.

WO 01/14961 A2



patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

— *Without international search report and to be republished upon receipt of that report.*

Title of Invention: System and Method for the Establishment and Utilization of Networked
Idle Computational Processing Power

Field of the Invention:

This invention relates generally to computer use over a network to form a distributed computing platform. More specifically the present invention is a system and method for use of idle time of computers that are connected to a network by users requiring significant computing power without the need for a large scale dedicated processing unit.

Background of the Invention:

When a computer is operating, but not actively performing computations for someone, it is said to be *idle*. Because of their incredible speed, modern computers are idle most of the time, not only when they are running screen savers, but *even when they are being used*. For instance, a fast typist working at top speed in a word processor uses only a fraction of the available computational capacity of a desktop PC. Although the time between a typist's keystrokes seems immeasurably small to humans, it is an eternity when measured in computer time - an eternity filled with unproductive *idle computation*. Because of this gross underutilization, it is estimated that well over 95% of the world's computational capacity presently goes to waste.

Summary of the Invention

The present invention is drawn to organization and use of the idle processing power of general-purpose computers to form a distributed computing platform.

It is therefore an objective of the present invention to provide large amounts of computational power to users without the users having to purchase a large computer for such purposes.

It is a further objective of the present invention to harness the idle computational power of many computers and make that power available to clients on an as needed basis.

1 It is yet another objective of the present invention to allow those who provide
2 computers to be used to have complete access to their computers whenever desired without
3 interruption of personal use.

4 It is still another objective of the present invention to provide computational power
5 without regard to any specific schedule of computer non-use.

6 It is yet another objective of the present invention to allow clients who require
7 computational power to specify the characteristics of the power they require.

8 It is still another objective of the present invention to charge clients that need
9 computational power based upon the power and other characteristics specified by the client.

10 It is a further objective of the present invention to permit the client to specify the
11 reliability of the computational power required by the client

12 It is yet another objective of the present invention to allow a provider of a computer to
13 be used in the present invention to specify the amount of RAM and disk space the provider
14 wished to devote to use by others.

15 It is still another objective of the present invention to push content of various types to
16 providers during the time that providers are supplying excess computational power to tasks of
17 the present invention.

18 It is yet another objective of the present invention to provide security to both the client
19 and the provider during the running of tasks on providers.

20 It is still another objective of the present invention to insure that a client cannot access
21 the computational resources of the provider beyond that which is required to run the clients
22 tasks.

23 It is a further objective of the present invention to permit clients to specify checkpoints
24 in the running of client tasks so that interrupted tasks can be resumed without significant the
25 loss of work.

26 It is still another objective of the present invention to charge clients for the use of
27 computational power in a flexible way based upon the amount of power to be used and the
28 reliability of the power to be used.

1 It is yet another objective of the present invention to allow clients to create a virtual
2 cluster of machines of client-definable computational power to run client tasks.

3 It is thus an overall goal of the present invention to utilize this heretofore underutilized
4 computer power in a novel way as a distributed computing platform to meet the needs of users
5 who require vast computing power but who may not have the financial wherewithal to
6 purchase or lease large mainframe computers or other supercomputing solutions.

7 Just as email service providers deliver email from one user to another, a business
8 running the system and method of the present invention will disseminate the tasks of a client's
9 distributed computer program to multiple providers for remote execution. As the remote
10 providers complete tasks, the providers will deliver results back to the originating client. As
11 part of the present invention, an intermediary server delivers *computation* from providers to
12 clients.

13 The Internet infrastructure to facilitate the distributed processing of the present
14 invention requires three components:

- 15 1) Client software applications for requesting and submitting distributed processing
16 requests and jobs. This is accomplished through a client API that allows client
17 jobs to be coded as independent tasks. The underlying task model, while complex,
18 is hidden from the user.;
- 19 2) Provider software called a compute engine (CE) for managing the launch and
20 execution of tasks delivered from the intermediary server. The CE runs
21 unobtrusively and securely on the provider's computer. It processes tasks when
22 the computer is idle and returns results to the server when the provider computer
23 connects to the network which, for example, and without limitation may be the
24 Internet; and
- 25 3) A centralized task server for exchanging tasks and results between participants.
26 The centralized task server (CTS) comprises a three-tiered architecture as more
27 fully set forth below and is redundant, fault tolerant and scalable as traffic and
28 clients increase. A task scheduler in the CTS matches power requested to
29 available provider computer resources.

1 For purposes of this application, the client is the individual or organization that is
2 requesting the data processing services of the present invention. The provider constitutes the
3 plurality of individuals or organizations having excess processor capacity that are offering
4 that capacity to the system of the present invention to process requests from clients.

5 The flow of the present invention is as follows:

- 6 1. Using the client software from their desktop, clients launch their distributed programs,
7 the tasks of which are transported to the centralized task server.
- 8 2. Client tasks are pooled at the centralized task server and allocated to provider
9 computers based on a scheduling algorithm that takes into account the characteristics
10 of the provider computer (e.g., processor speed, disk space, amount of RAM,
11 communication speed, percentage of time the provider computer is on-line; percentage
12 of time the provider computer is in use). The CE on the provider computer
13 periodically contacts the server of the present invention to retrieve task assignments or
14 other control messages. Such contact occurs periodically and opportunistically
15 (whenever a network connection is present, for example, when a provider computer
16 has dialed into the Internet) and not just when idle computational power is detected.
17 The invention also allows for an auto-dial capability where the provider computer can
18 specify intervals when the CE can connect to the server. Multiple tasks can be
19 assigned to a CE at a given time. This is done to increase the probability that the CE
20 will always have tasks to execute.
- 21 3. When a provider computer's CE detects that its host computer is idle, it executes tasks
22 that were retrieved from the centralized task server.
- 23 4. The results of completed tasks are saved locally. When a connection to the centralized
24 task server exists, statuses of current tasks, and results of previously executed tasks are
25 sent to the centralized server of the present invention.
- 26 5. Results are pooled by the centralized task server until the originating client can
27 connect to the server and retrieve them. Note that a special case of this is the case
28 when the client is connected as the job is running. In this case, the client sees the
29 results in "real time".

1 In this fashion, from a client's perspective, the present invention represents a gigantic virtual
2 computer or distributed computing platform, ideally suited for performing large-scale parallel
3 *distributed computation*, a function that was formerly the exclusive domain of traditional
4 "chips-and-solder" supercomputers. The present invention, however, will perform such
5 computations at a fraction of the cost and on an as-needed basis.

6 The business model associated with the present invention is to create a new market for
7 idle computation power. A business running the apparatus and method of the present
8 invention will purchase the idle computational power of millions of *providers* (individual
9 computer owners), and deliver that idle computational power as a single entity to *clients*
10 (technologists seeking ultra-high-performance computation).

11 Presently, idle computation has no value whatsoever because the market for
12 computation is inseparable from the market for computers. Users have no means of buying,
13 for example, larger amounts of times on disparate computers to execute a job when the user
14 needs it and not just when a specific computer is available.

15 With the present invention, however, immense value can be created by a
16 computational intermediary in the same way that financial engineers create value by
17 constructing derivative financial instruments. In effect, the present invention *strips* excess
18 computational capacity from provider computers, *bundles it* into quantities suitable for large-
19 scale parallel computations, and *delivers* it to clients on an as needed basis. Providers benefit
20 by selling a heretofore-wasted resource. Clients benefit because they can inexpensively
21 purchase just the quantity of computation needed, thereby avoiding the normally huge fixed
22 costs associated with purchasing a supercomputer.

23 The business model of the present invention creates and intermediates the new market
24 in idle computation.

25 Finally, an applications research and development group will add value by utilizing its
26 superior knowledge of distributed processing to solve high-payoff computational problems.
27 Such a team will spur demand in two ways: First, its purchases of idle computation will
28 stimulate market activity; second, the success of applications and research group will
29 demonstrate to prospective clients the power and cost-effectiveness of idle time processing.

1

2 **Brief Description of the Figures**

3 **Figure 1** illustrates an architecture of one embodiment of the present invention
4 **Figure 2** illustrates a flow of work over one embodiment of the present invention..
5 **Figure 3** illustrates one embodiment of the CE architecture of the present invention.
6 **Figures 4A and 4B** illustrate a normal execution scenario for tasks.
7 **Figure 5** illustrates an overall architecture of a server of the present invention.
8 **Figure 6** illustrates a determination of processing time of virtual nodes.
9 **Figure 7** illustrates communications paths of the present invention.
10 **Figure 8** illustrates expected time for completion of a task probability.

11

12 **Detailed Description of the Invention**

13 The present invention is a system and method for allowing those people who own
14 personal computers or larger computers, (hereinafter "providers") to offer unused and excess
15 computational power on their "provider computers" to people who have need for large
16 computational power for their computing tasks. These people are known as "clients."

17 Referring now to **Figure 1**, the architecture of the present invention is illustrated. The
18 present invention basically comprises three major components: a client application that has
19 been developed within a client development environment; a centralized task server (CTS) 26;
20 and a compute engine (CE).

21 In the context of this description a client application is an application program,
22 developed within the client development environment that is part of this invention, which
23 solves a problem utilizing the system that constitutes this invention. The typical problem such
24 an application solves is an extremely computation intensive problem that is amenable to
25 decomposition into a large number of independent tasks (this is not intended as a limitation).
26 The client development environment consists of a runtime API which defines how tasks
27 interact with their environment, a client API which allows a client application to manage the
28 execution of jobs consisting of a very large number of tasks, and a set of tools that support the
29 monitoring and control of very large jobs.

1 A centralized task server (CTS) 26 intermediates the needs of client applications and
2 the power of the provider computers 30-36 to meet the needs of clients for computational
3 power. Clients make requests via client computers 20-24 and by virtue of a client application
4 programming interface (Client API) that allows the client to specify the amount of computing
5 power desired, the individual rating of computers that will execute the client's job (e.g., 2
6 gigaflops or GF [Flops = Floating point operations per second]), how much memory and disk
7 space will be necessary and other parameters essential to running the client's job. The CTS
8 receives tasks from clients and assigns them to one or more providers based the characteristics
9 of the available providers and the job characteristics specified by the client. When providers
10 return results, the CTS forwards these back to the client application.

11 Each provider computer has a resident compute engine (CE) that receives and
12 executes the various tasks that are assigned to it by the CTS. The CE runs client tasks and
13 manages the use of idle computation power of the provider computer. The CE returns
14 intermediate and final results to the CTS that then forwards them back to the client
15 application.

16 Communication between the client application, CTS, and CE is governed by a
17 communication protocol unique to the present invention. At a high level of abstraction,
18 separate message and data channels exist. The message channel is a bi-directional channel
19 between clients and the server that is used to pass high-level messages that facilitate
20 distributed processing. The data channel is a channel used to pass large blocks of data called
21 data elements and executable elements. Executable elements are blocks code developed by
22 clients or a third-party. This code represents tasks or portions of tasks. Data elements are
23 blocks of data that are either input to or output of client tasks. High-level messages
24 transmitted on the message channel contain references to data and executable elements.
25 These items are separately transferred on the data channel on an as needed basis.

26 The details of the client development environment, CTS, CE, and protocol the
27 protocol that these components use to communicate will be covered in subsequent sections.

28 Referring first to **Figure 2** the general flow of the present invention is illustrated.
29 Clients 10, that is, those entities desiring execution of computationally intensive tasks, launch

1 client software that submits jobs consisting of multiple tasks. Although only one box is
2 shown, it is anticipated that there will be many clients that have such computationally
3 intensive jobs on or about the same time. Thus more than one client will be requesting such
4 services.

5 The centralized task server (CTS) takes all incoming tasks and pools them in some
6 priority order. The CTS validates each task (checks that all required executable and data
7 elements are in place) and assigns it to one or more CEs, based on information on available
8 provider computers and their characteristics. The software of the CTS will be able to size the
9 available computational power to the tasks that are pooled thus assigning the tasks the
10 appropriate amount of available computational power.

11 Periodically, and depending on when a network connection is present, the CE contacts
12 the CTS to see if there are any messages for it. These messages could be task assignments or
13 other control messages. The CE then retrieves tasks from the CTS so that, to the degree
14 possible, the tasks are readily available for execution whenever idle computing power is
15 available. When a CE receives a task assignment from the CTS, it downloads the data and
16 executable elements required to run that task. The task is then put in the "prepared" state.
17 Each CE constantly monitors its processor state and when the processor is idle it launches the
18 "next" prepared task. The task selected to run is a function of the arrival time of the task, the
19 task state, and the task priority. If a task is received that has a higher priority than the
20 currently running task, the current task is preempted (halted and results saved) and the new
21 task is executed.

22 Once tasks are executed, task results are returned to the CTS. The CTS journals all
23 work done by provider computer CE's on behalf of clients. This journal is then used as the
24 basis for adjustments to client and provider computer accounts on a periodic basis using
25 known accounting methods. These accounts are the basis for client billing and provider
26 payment.

27 As mentioned above, clients develop applications using the client development
28 environment that is part of this invention. The client development environment consists of a
29 runtime API which defines how tasks interact with their environment, a client API which

1 allows a client application to manage the execution of jobs consisting of a very large number
2 of tasks, and a set of tools that support the monitoring and control of very large jobs. These
3 are described in more detail below.

4 The runtime API defines the environment in which client task run. The runtime API
5 consists of a set of function calls that the client task will exercise during execution to interact
6 with its environment. This API contains calls that the task can use to:

- 7 • Obtain access to the initialization parameters for the task.
- 8 • Obtain access to data elements needed to perform its function.
- 9 • Perform a checkpoint so that if the task is interrupted, it can restart at the checkpoint
10 and therefore reduce the loss of work.
- 11 • Send interim and final results back to the server.

12 In addition, the API defines entry points that allow the environment to control the task.
13 These include:

- 14 • Start a task (note that restarting a task is a special case of starting a task).
- 15 • Stop (suspend) a task.
- 16 • Request a checkpoint.
- 17 • Request a status.
- 18 • Terminate a task.

19 This is not an exhaustive list of functions and is only illustrative of what can be
20 accomplished via the runtime API.

21 There can be multiple implementations of this API which allow the task to run in
22 multiple environments. The two most important implementations for the purpose of this
23 invention are the implementation in the CE and the run local implementation included in the
24 client development environment. The CE implementation allows the task to run in the CE on
25 a provider's box (regardless of platform). The run local implementation allows clients to
26 execute tasks on their local client machine. This enables debugging and testing before the job
27 is distributed. However, the client could elect to run an entire job (or a certain number or
28 percent of tasks) locally.

1 The client API is used to develop an application to control the execution of a job running
2 on the present invention. The client API broadly supports job definition, job launch, job
3 monitoring, and job control.

4 Job definition involves the creation of the job, setting job properties (e.g., requested
5 power, deadline, maximum cost, scheduling hints), creating its data/executable elements and
6 its tasks, and associating data/executable elements with tasks. The client has the option of
7 providing certain limits to the job being submitted. For example, a client can limit the job to a
8 certain dollar expenditure or amount power applied.

9 Job launch involves launching the tasks that constitute a job.

10 Job monitoring involves the selective monitoring of a job. The data available included
11 both status and statistics. The client API allows the client to attach to a job or a specific task
12 and monitor all or a subset of the results being returned by the job or task. Tasks can return
13 interim and final results via the runtime API. Certain default patterns are set by the system
14 concerning when interim results are to be obtained although the client can change these
15 parameters as well as explicitly request results for executing jobs.

16 Job control functions include starting, removing, stopping, and resuming jobs and tasks.

17 The client API of the invention supports jobs of arbitrary size. A "fire and forget"
18 event model allows clients to launch jobs much larger than the client machine would
19 otherwise be capable of tracking simultaneously. This is accomplished by programmatically
20 ensuring that clients monitor only task results, not the tasks themselves. This approach
21 serializes the receipt of task results such that a limited size computer can "drive" an arbitrarily
22 large job.

23 The invention also includes general-purpose tools, developed using the Client API, for
24 monitoring jobs. Jobs can be monitored whether they run locally (at the client's node) or
25 remotely (on a provider computer node). The client display is updated periodically (the
26 periodicity is configurable). The display for the client is designed to keep the client fully
27 informed at all times regarding the job submitted. As such it includes dollars spent, number
28 of tasks running, number of completed tasks, number of task that are executing, the number of
29 task delivered to the CTS, the total work accomplished, the amount of work done per task

1 (percentage), the work done on completed tasks, and the time spent per completed task.

2 These statistics are illustrative only. Other statistics may also be gathered as desired by the
3 client. The information that is monitored is displayed in both graphic and text form.

4 These tools can also be used to support other general-purpose functions such as
5 starting and stopping individual tasks or jobs, removing individual tasks or jobs, etc.

6 Referring to **Figure 3**, one embodiment of the CE architecture of the present invention
7 is illustrated. The CE is the software unit that resides on a provider computer and
8 accomplishes the timely and secure execution of client tasks. The CE **84** manages the
9 interface between the provider computer and the server of the present invention, monitors the
10 state of the provider computer's machine, schedules the execution of tasks on the provider
11 computer, and manages all communication with the server of the present invention. For
12 example, the CE **84** receives tasks from the CTS, posts results to the server, provides status of
13 tasks and receives messages to remove, suspend or resume tasks.

14 The CE **84** can be configured to execute tasks continuously, albeit at low priority (so
15 as not to interfere with a provider's normal use of their computer), or to start execution of
16 tasks when a provider's computer is idle (i.e., not being interactively used) and stop execution
17 immediately upon provider interaction with their computer. In this way, the CE is designed to
18 be unobtrusive to the provider.

19 CE **84** communications with the server and execution of tasks occur independently.
20 For example, even if a provider is actively using their computer, the CE communicates with
21 the server, provided a network connection is available (e.g., when a provider has dialed into
22 the Internet). In so doing, the CE receives elements that are placed in the task queue **86** to be
23 executed whenever possible. Thus, obtaining and staging of client tasks takes place at a time
24 independent from the execution of those tasks. Making CE communications independent
25 from CE task execution adds to the efficiency of the overall system.

26 The CE also includes a dialing scheduler that can be preset by the provider to dial out
27 to the server of the present invention at configurable intervals and only during preset periods.
28 Thus, for example, the provider can direct the dialer to dial out to the server every ten
29 minutes, but only between the hours of 2 a.m. to 4 a.m.

1 The entire CE resides on the provider computer, which as noted earlier, can be any
2 form of computer having excess capacity. For example, but not intended as a limitation,
3 provider computers that run the Windows® operating system, the Macintosh® operating
4 system, the Linux® operating systems and others are all candidates for use with the present
5 invention.

6 Various security measures are taken in executing tasks using the present invention.
7 The CE uses a virtual machine 78 to execute client tasks 82. Examples of virtual machines,
8 but not intended as a limitation, include the Java Virtual Machine and the C# Virtual
9 Machine. Throughout most of the remainder of the discussion we refer to the CE virtual
10 machine as the Java Virtual Machine or JVM 78.

11 Security restrictions on the virtual machine prevent client tasks from accessing
12 resources on the provider machine directly. For example, within the CE, the virtual machine
13 is altered such that only encrypted tasks provided by the server can be executed. This both
14 protects tasks when stored on disk and ensures that the CE executes only valid tasks. Further,
15 the system has the option of encrypting everything that is written to disk both in the provider
16 computer and on the server.

17 Referring to **Figures 4A and 4B**, a normal execution scenario for tasks is illustrated.

18 It is important to note that the CE 84 is a highly capable application. It has the ability to
19 run tasks constantly at low priority so that a provider's normal use of the computer is
20 unimpeded; alternatively, it can actuate in much the same way that a screen saver actuates,
21 executing tasks only when a provider is not using their computer. In practice, the CE will
22 exercise a small window that will be visually perceptible by the user of a particular computer
23 of the present invention. When the computer is idle for a configurable period of time, the CE
24 window will "pop up" with the cursor within the CE window. Alternatively, if the user so
25 desires, the application window can also simply be illustrated with an actual screensaver, but
26 the description below assumes the CE window is an application window.

27 The CE window activity is independent from the execution of tasks, but can be
28 synchronized with such so as to provide, among other things, status information about the
29 execution of the current task. Various buttons and areas within this CE window will be

1 accessible to the user such that accounts can be accessed, amount of time associated with use
2 of the users computer can be determined.

3 The graphic window of the CE is displayed for the provider when the CE detects the
4 host computer is idle. Its presence indicates to the provider that the system is processing, or is
5 ready to process tasks in the queue, and has a number of functions. The window can show the
6 state of processing, that is whether tasks are being processed at a particular moment or not.
7 Further, within the window are located various other buttons that may be clicked to determine
8 the status of the provider's account. In addition, advertising or other messages may be
9 presented in the CE window, which will allow a user to proceed to a particular website to
10 view further information concerning the messages. Thus, the CE window can operate in the
11 nature of a browser.

12 The size and placement of the CE window during "pop up" is configurable by
13 the user. When the CE window is displayed, the system seizes control of the cursor and
14 locates it at any configurable spot within the CE window. Thus the cursor can be placed in a
15 spot where messages are to be displayed or in a location to actuate any functionality that is
16 available to the provider while the CE is operating, such as the quick disappearance of the
17 display window as described above.

18 The user can also confine the CE to a system tray if desired. In this case, the icon or
19 panel in the tray can be used to indicate, among other things, engine status, the arrival of new
20 display content or upgrade packages, or Internet connectivity status. Alternatively, the CE
21 can be configured and executed as a "daemon" having only a command-line user interface.

22 The CE also supports the dynamic display of marketing content in the CE window to
23 the provider computer. The server notifies the CE when new content is available and where it
24 can be downloaded.

25 In the event that the user desires to use the computer for purposes other than to run
26 elements of the present invention, the user will simply type a key or cause the cursor to be
27 moved outside the CE window. As soon as this occurs, the CE will shut down all of the
28 processing of the present invention in an orderly fashion. Alternatively, the CE can be
29 configured to execute tasks constantly, albeit at low priority, in which case the CE window

1 can be made to disappear as described above, yet tasks will continue execution at lower
2 priority.

3 If the user would like to interact with the window, then a single click anywhere on the
4 window will cause it to enter a new state where it will not automatically disappear. This
5 contributes to the ease of use of the CE. This state is exited when the user closes the window
6 or after a period of inactivity.

7 As noted above, there are a series of execution steps when a CE is running on a user's
8 computer. For example, and referring to **Figures 4A and 4B** at the outset, the CE obtains a
9 task message from the server **90**. The executable elements from the element server are
10 written to the disk of the computer. However, before writing any items to disk, the disk
11 management logic within the CE ensures that any such storage will not cause the CE to
12 exceed its storage space budget **92**.

13 The CE then writes out the task policy file that establishes the permissions for the
14 Runtime **80**, and the task to be run **82**. This task policy ensures that the client task can neither
15 access the provider computer's disk or the network. The runtime policy gives the runtime
16 module **80** access to data elements need to run the task (if any) and authorizes the runtime
17 module **80** to connect back to the Core Engine. The CE then launches the virtual machine
18 (JVM) **96** and executes the Runtime **80**. The path of the policy file, the security manager, the
19 maximum size to be used based upon the CE's memory budget, and the class path including
20 all executable elements, are passed to the JVM **96**. The class name for the task and the CE
21 communications port are also passed as parameters to the runtime module **80**.

22 The JVM then launches the Runtime main module **98**. This runs in a protection
23 domain that allows the Runtime to open a network connection to the CE and to read
24 executable and data elements from the disk. The JVM opens a server socket, connects to the
25 CE port passed in on the command line, and sends a message containing its server port **100**.

26 The CE then connects to the JVM's server port. This results in two connections
27 between the CE and the JVM, i.e., a CE to runtime channel, and a runtime to CE channel **102**.
28 The runtime requests the parameters for the task over the runtime to CE channel **104** and the
29 CE returns those parameters in response over the CE to runtime channel.

1 The runtime then requests the data elements for the task **106** over the runtime to CE
2 channel, and the CE returns them in response over the CE to runtime channel.

3 The runtime then establishes a task context that the task uses to communicate with the
4 runtime system **108**.

5 The runtime establishes and begins the client task **110**. The executable elements for
6 this task having been provided to the runtime via the class path provided for the JVM when
7 the JVM was started. The client task then runs in a separate protection domain that prevents it
8 from accessing the network or the disk. This protection domain is established through the
9 permissions granted to the task in the task policy file.

10 The client task interacts with the runtime through the task context object established
11 for it by the runtime **112**. The runtime calls run with the client tasks privileges when possible.
12 If the runtime needs to perform an operation requiring additional privileges, i.e., accessing a
13 particular data element, the runtime executes a privilege block that takes on the minimum
14 privilege for the minimum amount of time required to complete the operation.

15 The client interacts with the runtime to get task parameters and to access task elements
16 **114**. As an optional step, when the client task can send out temporary results using the set
17 status call to the runtime **116**, these results are passed to the CE over the runtime to CE
18 channel. The CE then sends these interim results in the form of a task status to the server.

19 As an additional option, when the client task can create a task checkpoint using the
20 checkpoint call to the runtime **118**, the checkpoint is passed to the CE over the runtime to CE
21 channel. The CE serializes the checkpoint to disk so it is available in case a task restart is
22 required.

23 When finished, the client task does a set status runtime call **120** with the final results
24 and exits. The results are passed through the CE over the runtime to CE channel. The CE
25 then sends these results in the form of a task status message to the server. After the CE sends
26 a final task status noting that the task is complete **120**, the space that is consumed by the task
27 definition and related data can be reclaimed for use in running other tasks.

28 Once the task completes **122**, the runtime closes down its connections and exits.

1 As noted above, the CE is a very flexible program that does not impair the ordinary
2 use of the computer in any way. For example, when the user moves the mouse cursor outside
3 of the CE window, the CE sends a shut down request to the Runtime. If the JVM does not
4 shut down in seconds, the JVM is automatically stopped. When the Runtime receives a shut
5 down request, it calls the clients task stop method. The client is allowed to send a set status
6 and create a checkpoint before exiting. This reduces the loss of work and allows the program
7 to restart at the checkpoint when processing resumes later. If the client task supports
8 checkpointing, the runtime sends the checkpoint back to the CE on the runtime to CE channel.
9 If a checkpoint is received, the CE writes it to disk. Once the client task exits, the runtime
10 closes its connections and exits.

11 If the CE has been configured to execute tasks intermittently based on user interaction
12 (versus continuously), then when the CE detects that the PC has been idle for a configurable
13 period of time, and thus is available to run a task, it checks its queue for a paused task, and
14 restarts the task from the last checkpoint which was previously stored on disk, using the same
15 mechanism used to start the task at the outset. If there was no paused task, the next task in the
16 queue is run.

17 It should be noted that the CE functionality, while used in the present invention, could
18 also be used in other fashions. This would, for example, allow particular programs to be run
19 within a given office, thereby utilizing the full power that is available on all PCs within an
20 office location. Again, processing occurs whenever the computer is idle and ceases whenever
21 the individual PC user desires to use the PC for specific purposes.

22 The centralized task server (CTS) intermediates tasks submitted by a client application
23 time and the CEs on providers' computers that have excess computational capacity to run
24 those tasks. The CTS registers clients, accepts tasks for execution, distributes those tasks to
25 computers, receives results back from the provider computers, conveys results to the client
26 computers, charges clients for the computational work performed and pays providers for the
27 use of their computers.

1 One of the most critical functions of the CTS is task scheduling. Part of the
2 scheduling process is to provide clients of the present invention, who desire to obtain the
3 aggregate computational power, with the desired amount of processing power on demand.

4 The server of the present invention loads tasks onto a single provider computer based
5 in such a way as to match computer capability with task requirements. As tasks are
6 completed, and through the communications channel between the provider computer and the
7 server, additional tasks are fed to the provider computer for execution. All new tasks are
8 placed in a task execution queue on the provider's computer for subsequent execution. To
9 provide maximal control of the execution of tasks on a provider's computer, the CE supports
10 the concept of task priorities and task preemption (tasks can preempt executing tasks of lower
11 priority). These features allow the server to more effectively utilize available provider
12 computer resources.

13 In order for the server of the present invention to correctly characterize the provider's
14 computer, the server also collects usage and performance histories (profiles) from the provider
15 computer concerning a wide variety of parameters. For example, the server collects Internet
16 connectivity patterns and processor availability patterns during the course of any 24-hour
17 period, weekly period, monthly period, etc. The processing power of the processor itself is
18 also important since certain clients may request higher or lower processing power depending
19 on relative cost and timeliness tradeoffs.

20 To fulfill its function, the task scheduler must fulfill the following requirements: The
21 task scheduler gets jobs and associated tasks from clients. It must then take those tasks and
22 distribute them to various provider computers for execution. It must also record the fact that
23 it has assigned a particular task to a particular provider computer.

24 The task scheduler passes both interim and final results to the client and records the
25 amount of work performed to an on-line ledger. It then credits a provider's account with the
26 amount of time that the provider computer actually performed on the specific task. This
27 allows the provider to be paid by the computational work that it has contributed to a particular
28 job. (By definition, time multiplied by the power of the provider computer equals work
29 performed.)

1 The task scheduler also logs completed work in a client's account noting the amount
2 of work performed so that an appropriate billing can be passed on to the general ledger
3 module for billing to the client.

4 Cancellation of tasks or jobs (collections of tasks) from a client is also processed by
5 the task scheduler. It accepts any job cancellation from the client and insures that appropriate
6 communication is sent throughout the system of the present invention canceling a particular
7 job.

8 When a client job is completed, the task scheduler automatically provides e-mail or
9 other automated notification to the client that the job is complete.

10 Results that are received by the task scheduler are archived if the client has not
11 retrieved those results. From time to time the task scheduler deletes unclaimed information
12 from the data archive.

13 Since clients pay for the computational power they use, the task scheduler periodically
14 checks the on-line ledger to determine if a client has exceeded an account limit. If so, the task
15 scheduler notifies the client and, if necessary, cancels all jobs for the client until such time as
16 the account is brought current.

17 In the event that a particular provider computer that is working on a particular task
18 does not respond for a pre-set period of time, the task scheduler moves the task to another
19 provider computer. Thus, it must be able to detect unresponsive CEs on provider computers.
20 Depending upon the time constraints and priority of tasks involved, the task scheduler also
21 can elect, through its software, to launch redundant tasks to different provider computers to
22 insure timely completion of any particular task.

23 The CTS is also tasked with delivering the appropriate amount of compute power for a
24 client task, not more or less, for a given client. Since charges for the client are made in part
25 based upon the compute power delivered, this assembly of provider computers and delivery of
26 computational power must be accurate.

27 In order to assure responsiveness to clients, the task scheduler also balances data
28 processing loads among the various provider computers based upon the capabilities of the

1 provider computer and the bandwidth that is available to provide elements and tasks to
2 provider computers.

3 The CTS collects statistics regarding provider computer capabilities. This is critical
4 for both assigning payment to a provider and a debit notice to the client. It is also important
5 that the CTS know at all times the capability of the provider computers.

6 As noted above, the CTS may send the same tasks to more than one provider
7 computer. This helps the CTS be assured that the computations that are performed by the
8 assigned provider computers are correct. Since results from provider computers should be the
9 same for the same task (unless randomization is integral to the calculation of the task), any
10 differences in response are noted by the CTS and resolved. This assures the client that
11 answers are accurate.

12 The CTS takes explicit measures to protect the intellectual property inherent in a task.
13 To accomplish this, the CTS obfuscates executable elements and data elements to remove
14 contextual clues to the code's intended purpose. Furthermore, the CTS also removes from
15 tasks all traces of client identity. Thus, unless specifically intended otherwise by the CTS, a
16 provider has no means of knowing the clients for whom elements are executed nor their
17 domain of inquiry.

18 To characterize and assess the capabilities of provider computers, in order to perform
19 task scheduling effectively, the CTS executes certain benchmark tasks on provider computers.
20 Also, various test tasks with known results are sent to the provider computer to ensure that
21 results returned by the provider computer are accurate. Unexpected results from a provider
22 can be used to identify attempts to submit fraudulent results or otherwise tamper with the
23 normal execution of the system of the present invention.

24 In operation, the present invention allows the client to specify job and task parameters
25 via the client interface. Using the client interface, the Client selects desired attributes of
26 nodes in a virtual cluster. These attributes include, but are not limited to, CPU speed (in GF),
27 memory, available disk space, intermittent or constant connectivity, bandwidth (in Kbps),
28 number of nodes, and reliability (as measured by consistency of results). The client also

1 provides information on his job: estimated work per task (in GF-Hours), incoming data (in
2 MB), outgoing data (in MB).

3 To deliver a virtual cluster with the desired attributes, the CTS may construct virtual
4 nodes (one or more nodes operating on the same task so as to improve the expected time-to-
5 completion of a given task). Virtual nodes are discussed in more detail elsewhere in the
6 document.

7 After the client defines the tasks to be executed, they are provided to the central server
8 of the present invention for "launching," i.e., the placement of the task in a queue for
9 scheduling and ultimate distribution to provider computers. For billing and compensation
10 purposes, a task is "launched" when it is scheduled. In this way, the latency from a task sitting
11 in queue is attributed to the provider computer and so considered part of the provider
12 computer's intermittent "unavailability".

13 The system of the present invention records the time at which each task is launched to
14 the first provider computer in a virtual node. The system also records the time at which the
15 first completed task is returned from a virtual node. At any point in time, the amount of power
16 running through the virtual cluster is the sum of the powers of the virtual nodes running.

17 The system constantly monitors virtual nodes for possible failure. Failure is defined
18 as an absence of communication from a configurable number of provider computers in a
19 virtual node for a configurable amount of time. When a virtual node fails, the task assigned to
20 it is re-launched on another (perhaps faster) virtual node.

21 Below we describe a novel pricing model and task scheduling model for use with the
22 system of the present invention, but neither is intended as a limitation.

23 At any point in time, the client can be billed. The client charge is the price per hour
24 (based on power running through the cluster) for each hour (or portion thereof) since the last
25 billing. For example, and referring to **figure 6**, the determination of running time on
26 providers' computers is illustrated. In the example, each virtual node in the cluster operates at
27 0.1 GF. The client is charged an hourly rate for running tasks based upon the power, P ,
28 running through the client's account at each 15-minute increment. The price per hour for
29 computation is:

$$\text{Price per hour for computation} = 4.5 + 0.296P^{1.155}$$

where 4.5, 0.296 and 1.155 are configurable constants.

The client is charged an additional amount per GB for data transferred through each of the four communications paths, **150**, **152**, **154**, and **156** of the present invention as illustrated in **Figure 7**. The communications paths **150**, **152**, and **154** represent duplicated data and are a function of the degree of reliability the client achieves. Of the dashed arrows, only one will be executed, and so the client will be charged for only one of these paths.

The client is also charged for data transfer during the course of executing the particular tasks of the client. For example, let d_i be the quantity (in GB) of data transferred from the client to the server. Let d_o be the quantity of data transferred from the server to the client. Given that some of the data will be duplicated for processing by virtual nodes comprised of more than one provider computer, the client's total data transfer charge is given by the following:

$$\text{Total Cost of Data Transfer} = g(d_i + d_o) \left(\frac{\ln(1-r)}{\ln(1-\pi)} + 1 \right)$$

where g is the transfer cost per leg per GB, r is the requested node reliability, and π is the average provider computer reliability.

Requested node reliability r can be defined for a virtual node of CPU speed C and bandwidth B as the probability of the virtual node completing a task in no more time than the time it would take a fully idle, constantly connected computer with CPU speed C and bandwidth B to complete the task.

Average provider computer reliability π is defined as the probability of the provider's frequency of contact not diminishing, in the near term, from the provider's historic average.

The client is also charged an hourly premium for requesting reliability greater than a set baseline reliability. The charge per hour is:

$$\text{Price per hour for premium reliability} = kCN \left(\frac{\ln(1-r)}{\ln(1-R)} - 1 \right)$$

where k is the baseline price per MHz-hour (provider computer price per hour for the ideal machine), C is the CPU speed requested of the virtual nodes, N is the number of virtual nodes running, r is the requested reliability, and R is the baseline reliability.

Provider Payments and Task Scheduling (Simplified Version)

As clients are charged for provider computers' time, so must the providers be paid for the computer time they provide. For each provider computer, a CPU availability rating, a_c , is constructed along with a bandwidth availability rating, a_b . These ratings measure the average availability of the provider computers' CPU and bandwidth over time. Unlike in the full version of the task scheduling model (described below), these ratings are single values versus vectors.

To measure availability, the system of the present invention records transitions of CPU and bandwidth (between "available" and "unavailable") over time. Defining each minute of availability as having value 1 and each minute of unavailability as having value 0, the system takes the moving average of these values over the past seven days to obtain average availabilities for CPU and bandwidth.

For each provider computer, the system next measures CPU speed, C , and bandwidth, B (through benchmark tasks). These measures are periodically updated to ensure that the system has knowledge of the full availability of the provider computer's system.

The client may provide estimates of the work w , data in d_i , and data out, d_o associated with the task. If the client does not specify these measures, the system uses the averages for a sampling of tasks.

The system then computes the expected time-to-completion for the task as using the following formula:

$$\text{Expected Time-to-Completion} = T_E = \frac{w}{Ca_c} + \frac{d_i + d_o}{Ba_b}$$

As noted earlier the client specifies the CPU speed, C_v , bandwidth, B_v , and reliability, r , of a virtual node. If reliability is not specified, the system of the present sets a default value for the baseline reliability R . The client may select a greater, but not lesser, reliability.

The system then computes the implied "requested" time-to-completion for the task as:

$$\text{Requested Time-to-Completion} = T_R = \frac{w}{C_v} + \frac{d_i + d_o}{B_v}$$

The subset of provider computers that satisfy the following criteria are considered "candidate provider computers":

- The provider computer is not working on a paying job.
- When the client specifies w , d_i , and d_o , the provider computer has (approximately) $T_E = T_R$.
- When the client does not specify w , d_i , and d_o , the provider computer has (approximately) $a_c C = C_v$ and $a_b B = B_v$.

Taking provider computers' CPU's and CPU availabilities as fixed, candidate provider computers have expected bandwidths (bandwidth multiplied by bandwidth availability) that satisfy the following equation:

$$Ba_b = Ca_c \frac{(d_i + d_o)C_v B_v}{wB_v(Ca_c - C_v) + (d_i + d_o)Ca_c C_v}$$

For each candidate provider computer, the system calculates the weighted average availability, π , as:

$$\pi = \frac{a_c \frac{\bar{w}}{C} + a_b \frac{\overline{d_i + d_o}}{B}}{\frac{\bar{w}}{C} + \frac{\overline{d_i + d_o}}{B}}$$

where the bars indicate the average over a sampling of tasks.

Selecting randomly from the candidate provider computers, the system arbitrarily groups the provider computers into virtual nodes such that the probability of the virtual node returning the task in, at most, T_R equals r . Specifically, the system randomly picks provider computers from the subset and groups them into virtual nodes such that, for reliability r , and N nodes, each with weighted availability π_i ,

$$r = 1 - \prod_{i=1}^N (1 - \pi_i)$$

It is important to note that the value of N that satisfies the equation may be non-integer. In this case, the server will launch a different number of duplications of each task such that the average number of duplications per task equals N .

The system launches the task simultaneously on all the provider computers that comprise a virtual node. As noted earlier the task is “launched” when it is scheduled. When the first provider computer returns the completed task or when the client cancels the task, all duplicate tasks on remaining provider computers are cancelled. We define “cancelled” as the placing of a “remove task” command in the queue. A particular provider computer’s task is considered concluded when the first of the following occurs:

- The provider computer receives notice that the task has been cancelled.
- L hours pass since the task was cancelled.

For each provider computer, the system computes the time that elapsed between the launching of the task and the conclusion of the task. The amount of work the provider computer completed is estimated as the time that elapsed multiplied by the provider

computer's power rating. Note that, effectively, the provider is being paid for time, not work. However, because the provider computer's power rating is a function of the provider computer's availability, the rate of payment the provider receives is a function of the average work per unit time the provider computer completes over time.

Providers are paid a fixed amount per unit time. The system established a target annual payment per MHz-hour, k , and a baseline reliability, R . A provider computer with a weighted average availability of π and CPU of C is paid $\frac{kC \ln(1-\pi)}{\ln(1-R)}$ per hour. This assumes that the provider computer is working over the whole interval from launch to conclusion of the task.

Provider Payments and Task Scheduling (Full Version)

As part of establishing a given provider computer as a viable source of computing power for the system, for each provider computer, the system constructs a CPU signature and a bandwidth signature. These signatures show the probabilities of the provider computer's CPU and bandwidth being available for a fixed time interval over time. For example:

Beginning of Time Interval	Pr(CPU available)	Pr(Bandwidth available)
1	30%	0%
2	40%	10%
3	50%	20%
4	40%	30%

The number of intervals should be enough to cover a reasonable circadian "cycle" (e.g. 1 week). The signature should be updated periodically.

For each provider computer, the system measures CPU speed by sending to the provider computer certain benchmark tasks that have known response times. The bandwidth is also measured. As noted earlier, these measures are periodically updated.

1 The system then combines the signatures with CPU speed and bandwidth to construct
 2 a profile matrix. The profile matrix shows the probabilities of the provider computer
 3 completing a task of given work, input data, output data, and time of launch within certain
 4 time lengths. Example: The tree below shows all possible paths of success and failure of a
 5 provider computer working toward completion of a task. Each step down in the tree represents
 6 a time increment. Left branches represent failures to work. Right branches represent successes
 7 toward work. Probabilities are associated with each branch based on the CPU and bandwidth
 8 signatures. The expected time-to-completion is the sum of the products of the probabilities of
 9 reaching each terminal node and the times required to reach the nodes. This concept is
 10 illustrated in **Figure 8**.

11 The client specifies the CPU speed and bandwidth of a virtual node. The system finds
 12 all provider computers that have an expected time-to-completion equal to the time-to-
 13 completion for a fully available node with the specified CPU speed and bandwidth. Call this
 14 time-to-completion the “desired time-to-completion”.

15 The system of the present invention arbitrarily groups these provider computers into
 16 virtual nodes such that the probability of the virtual node returning the task in the desired
 17 time-to-completion equals some fixed value.

18 The system then launches the task simultaneously on all the provider computers
 19 comprising the virtual node. When the first provider computer returns the completed task, the
 20 duplicate tasks on the remaining provider computers are cancelled.

21 The system computes the work each provider computer performed from the time the
 22 tasks were launched until the first task was returned. We calculate work as the number of
 23 hours the provider computer worked multiplied by the provider computer’s power rating.

24 Providers are paid a fixed amount per unit work. Payment is made on a per GHz-hour
 25 rate, T , and a target reliability r . A provider computer with an average (based on average task
 26 attributes and over all time intervals) availability of π is paid $\frac{T \ln(1-\pi)}{\ln(1-r)}$ per GHz-hour.

27 Referring now to **Figure 5**, the overall architecture of the server of the present
 28 invention is illustrated. It should be noted at the outset that while single modules may be

described, multiple instances of each such module may, in fact, be present in the architecture in order to allow for scalability to larger operations, to increase reliability so that no single point of failure exists in the system, or to have multiple instances of the entire server located at different physical locations to allow for redundancy and insurance against failure.

The server comprises a registration servlet **50**, which accepts registration requests over the network **28** from clients and providers who wish to participate on the system. Registration information is sent by the registration servlet to the registration manager **52** for processing.

The registration manager **52** receives appropriate information from the registration servlet **50** and assigns the client/provider computer to a database. The registration manager then generates a certificate and returns the certificate in the response to the client/provider computer. The provider computer certificate includes the provider public key, the provider-id (which is a hash of the provider public key), and the provider database ID. The client certificate comprises the client/user public key, which then acts as the client/user ID, the client ID, and the provider computer database ID.

The provider servlet **56** accepts various messages from providers as will be discussed more fully below, and routes them to the appropriate provider manager **58**. Further, the provider servlet **56** returns responses from the provider manager to the various provider over the network **28**. The web server of the present invention encrypts and decrypts requests, authenticates requests, routes them to the appropriate servlet, encrypts and returns responses. The provider servlet accepts GET and POST messages from providers. (These messages will be more fully explained below).

Provider manager **58** manages the interaction of the system of the present invention with the various providers. The provider manager receives and routes requests from other internal server components such as the client manager **68**, the task scheduler **64**, and the database server **54**. All administrative information and processing that relates to providers is managed by the provider manager **58**.

The element servlet **60** is used as the point where data elements and code are exchanged between clients and providers over the network. Elements are executables, or

1 portions of executables that are to be run on provider processors. Elements are uploaded to
2 the element servlet that then stores them on the file server. When a CE receives a task
3 assignment from the CTS, the CE downloads the required elements for that task from the
4 element servlet. File server 62 stores data and executable elements that are to be exchanged
5 between clients and providers. Client servlet 66 accepts various messages from clients and
6 routes them to the appropriate client manager 68. Client servlet 66 receives responses from
7 client manager 68 and returns the response to the client. Since responses to and from clients
8 proceed over network 28, the client servlet is invoked by a web server that is responsible for
9 encryption, decryption, and authentication.

10 The client servlet also manages sessions with a client. This is accomplished by the
11 client servlet assigning a cookie to the particular client for the particular session. If a request
12 is received from a registered client and the request comes without a cookie, a new session
13 cookie is created and sent to the client by the client servlet. In this fashion, individual
14 sessions can be identified and managed.

15 Client manager 68 is responsible for managing requests to and from clients. Client
16 manager 68 maintains a valid message channel for each client session. When the channel is
17 established, the client servlet accepts GET and POST requests and forwards them to the
18 appropriate CM. The client manager uses business routines within the client manager to
19 handle POST requests. Depending on the type of requests, such processing may involve
20 interaction with the database server 54, the task scheduler 64, and/or the provider manager 58.
21 A GET request is handled by the client manager 68 by returning the appropriate message from
22 the client output queue (explained below).

23 Task scheduler 64 manages the assignment of tasks to providers. The task scheduler is
24 responsible for tracking the states of the various providers, that is, whether the providers are
25 currently processing tasks, whether the providers are available, and, without limitation,
26 whether the provider has the appropriate processing power to process the specific tasks that
27 may be necessary. The task scheduler 64 receives messages from client manager 68 and
28 provider manager 58 regarding any changes in the status of clients and providers. When tasks
29 are assigned to specific providers, the client manager keeps a record of such task assignments

1 in the database 54. Various algorithms present in the task scheduler ensure that appropriate
2 tasks are assigned to the appropriate providers given priority of tasking and other prioritizing
3 data.

4 Database server 54 stores information on clients and providers. To accomplish this
5 task, the database server 54 comprises an account database relating to providers and revenue
6 accumulated by providers, a client database relating to the identity and contact information for
7 clients, job and task information, and a provider database relating to providers and
8 identification and contact information for providers. Such information, while shown as being
9 located on a single database server 54 can also be located on multiple servers to so that the
10 system can be scaled up to handle many more users as well as ensure privacy and security of
11 the information.

12 The web manager 70 intermediates access between the account section of the web site
13 and the database. The web manager 70 supports the web based registration, update of
14 registration and account information, and retrieval of account-specific and system-wide
15 statistics.

16 Monitoring and control module 72 comprises a set of utilities that support the overall
17 management of the system of the present invention. The monitoring and control functionality
18 72 provides graphical user interfaces for system operators to monitor the health of the overall
19 system. The monitoring and control functionality 72 polls all servers for health information in
20 the form of name-value pairs that contains status information and related attributes to all of
21 the servers. In this fashion, the monitoring and control 72 can keep track of the health of the
22 entire system. The monitoring and control module can also respond to asynchronous
23 notification of system events from servers.

24 Log server 74 contributes to the health of the overall system by collecting debug and
25 error log data. Any errors that are encountered throughout the system whether they are in
26 clients, providers, or any of the servers of the system are collected in the log server for
27 subsequent analysis and debugging where necessary.

28 As noted above, while single instances of the various modules have been indicated in
29 **Figure 5**, this architecture is not meant as a limitation. Additional instances of any given

1 component may be added to response to increases in demand from clients or increases in the
2 inventory of providers. For example, database server 54 comprises information on clients and
3 providers and is partitioned into two databases, one for a client database, and one for a
4 provider database. These databases can run on a single server or on dedicated servers. As the
5 number of clients and providers increases, additional databases can be added.

6 Client servlet 66 and provider servlet 56 are stateless, therefore allowing additional
7 components to be easily added to increase capacity. When a client request comes in over
8 network 28, the requests are allocated among available client servlets to increase the response.
9 The same holds true for the provider servlet 60, thus allowing responses to and from providers
10 to be rapidly responded to.

11 Element servlet 60 is also stateless, thereby allowing additional components to be
12 easily added to increase capacity as well. Thus, a large number of element servlets may
13 potentially exist which store elements on a given file server. Clients and providers allocate
14 their requests and elements across the element servlets.

15 File server 62 will store large amounts of client data. Additional servers can be added
16 to the system to increase this capacity. Thus, it is within the scope of the present invention to
17 have multiple file servers 62.

18 Provider manager 58 and client manager 68 can also be augmented by additional client
19 and provider managers respectively. Each client is assigned to a specific client manager.
20 This is accomplished by assigning clients to groups, with the group subsequently being
21 assigned to a client manager. For example, client manager 1 may handle groups 0-49, and
22 client manager 2 may handle groups 50-99. When a new client manager is added, clients are
23 reallocated among the various client managers in order to ensure the most efficient flow of
24 messages among the elements of the system.

25 Task scheduler 64 can quickly become a bottleneck when multiple clients and
26 providers begin to grow. If this is the case, additional task schedulers can be added to manage
27 a subset of tasks flowing to and from a provider. Again, tasks and providers can be assigned
28 to groups with groups being subsequently assigned to specific task schedulers. Again, as task
29 schedulers are added, a real allocation of tasks and providers among the various schedulers is

1 accomplished in order to even the workflow and ensure the most efficient functioning of the
2 system of the present invention.

3 In a similar fashion, web manager 70 can be augmented by additional instances of web
4 managers with requests being allocated among the various web managers. The monitoring
5 and control module 72 can also be represented by multiple instances of the monitoring and
6 control function. There can also be multiple instances of the log server 74. However, there
7 must be close coordination among the various log servers so that any trends and errors can be
8 detected and corrected.

9 Registration servlet 50 can also be augmented by additional components. Clients and
10 providers are assigned across the various instances of registration servlets.

12 The task scheduler 64 maintains an in-memory model of the state of all of the
13 providers (i.e. the provider profiles) and any tasks that it manages. Information to establish
14 the task scheduler is initially retrieved from database 54. It is thereafter updated as messages
15 are received from the provider manager regarding status of efforts of the various providers,
16 and from the client manger, which forwards relevant messages from clients to the task
17 scheduler.

18 The task scheduler also requests additional information about providers by cuing a get
19 task message (for retrieving task status) or a get cache contents for obtaining the cache
20 contents of the server for the provider. The task scheduler comprises various scheduling
21 algorithms. This task scheduler makes initial assignments of tasks to providers and migrates
22 tasks from one provider to another if necessary. When the task scheduler schedules a new
23 task, it records the assignment of the task and the provider to which it is assigned in the task
24 database 54. It further queues a task message to the provider. When migrating a task, the task
25 scheduler removes the task from the current provider by cuing a remove task message,
26 assigning the partially completed task to a new provider via a task message. All such
27 operations are reported in the task database.

28 The element servlet is the entity that stores data elements that are to be assigned to
29 providers for processing. Clients upload their data elements to the element servlet 60. The
30 element servlet authenticates the client as one that is permitted to store elements in the file

1 server 62. If the client is authenticated, the element servlet 60 stores the object in a physical
2 location on the file server 62 and notifies the client manager 68 that the object has been
3 uploaded. The client manager 68 in the client database 54 then records this information.

4 Client computers download elements to the element manager. The element manager
5 60 registers the fact that this has occurred to the client manager 68.

6 The notification that is provided by the element servlet to the client manager 68
7 includes the client ID, a locator for the data element, and the physical location of the data
8 element.

9 When a provider needs an element, it makes a request to the element servlet 60 over
10 the Internet 28. The element servlet 60 authenticates the provider, decrypts the ticket that the
11 provider offers to the element servlet, compares the ID of the provider making the request
12 with the provider ID and the ticket. If the ID's match and the ticket has not expired, the
13 element servlet 60 looks up the element on the file server 62 and returns it to the provider for
14 execution.

15 When the client terminates a particular job, all of the elements associated with that job
16 are marked as deleted. Deleted elements are removed from the file server 62 and from the
17 database 54.

18 Communication within the system of the present invention among the CE, the CTS,
19 and the client is governed by a communication protocol unique to the present invention. At
20 the high level, separate message and data channels exist. Messages between the CE and the
21 server occur over the message channel while executable and data elements go over the data
22 channel. Executable and data elements are referenced from messages passed over the message
23 channel.

24 The message channel is implemented above a low-level protocol that is part of this
25 invention. The low-level protocol supports the secure and reliable transmission of messages
26 between system components (CE, client, CTS).

27 The protocol is secure because it is built on Secure HTTP (HTTP over SSL). The
28 invention uses SSL to both provide confidentiality (encryption) of transmitted data and to
29 authenticate the parties at each end of the connection. Providers and client applications can

1 only access the server if they have a legitimate certificate. In fact, the server uses a hash of the
2 client's public key as a client identifier.

3 The protocol is reliable because it incorporates sequence numbers, acknowledgement,
4 and retransmission. Combined, these protect against dropped messages, out-of-order
5 processing of messages, and the processing of duplicate messages.

6 It is worth noting, the low level protocol provides the abstraction of a bi-directional
7 communication channel between components even though it is implemented on top of HTTP
8 that has a request-response model. A bi-directional model is simulated by having the client
9 periodically "poll" the server for messages. This poll is done on an exponential decay to limit
10 the amount of traffic generated by polling. If the CTS does not respond to a particular GET
11 message, perhaps because of heavy CTS load, subsequent GET messages are sent ever more
12 infrequently, according to an exponential decay, until some limit periodicity is achieved. In
13 this manner, heavy CTS can be automatically mitigated through less frequent CE GET
14 requests.

15 The maximum polling rate is present each time the client and server exchange a high
16 level message. The two basic operations in the low level protocol are the "GET" and the
17 "POST", both of which are initiated by the client. The POST operation sends a new message
18 (or set of messages) to the server. The GET is a poll operation to get any message (or set of
19 messages) that the server may have enqueued for the client. Since the low-level protocol is
20 based on HTTP it takes advantage of that protocol's features for multipart messages and
21 different message encodings (e.g., for compression). Further, multipart messages can be sent
22 in order to maximize communication efficiency. All messages are also subject to data
23 compression to limit the volume of data being transmitted and to conserve bandwidth.

24 The low-level protocol also supports the concept of session. Sessions are implemented
25 using cookies and provide for the resetting of a communication channel in the event that
26 synchronization is lost.

27 The following is a listing of high level message types used by the present invention in
28 its communications protocol. This list is illustrative in nature. Other message types can be
29 added to the present invention as the need arises. As noted above, a series of messages are

1 transmitted between the system and various elements. These messages are characterized as
2 Client-to-Server messages, Server-to-Provider messages, Provider-to-Server messages and
3 Server-to-Client messages.

4

5 Client-to-Server Messages

6 GetConfig:

7 This message is sent when a particular client wants to reestablish its
8 configuration. A GetConfig message causes the CTS to send a Config
9 message to the client.

10 Task:

11 The client uses this message to send the server a task to be executed.

12 ClientProfile:

13 The client sends this message on startup and whenever client settings change.

14 CreateJob:

15 Create a new job.

16 Stop:

17 Stop execution of a job or task.

18 Resume:

19 Resume execution of a stopped job or task.

20 RemoveJob:

21 Remove (cancel) a job.

22 GetJobProperties:

23 Requests a JobProperties message.

24 GetContents:

25 Requests a Contents message.

26 GetExecutableElement:

27 Request that an ExecutableElement message be sent describing a particular
28 executable element.

29 GetDataElement:

- 1 Request that a DataElement message be sent describing a particular data
2 element
- 3 RemoveExecutableElement:
4 Request the removal of the named executable element.
- 5 RemoveDataElement:
6 Request the removal of the named data element.
- 7
- 8 Attach:
9 Attach to a particular task or job. The attach message is used to subscribe to
10 receive status updates on the identified task or job.
- 11 GetTask:
12 Request that the identified task be checkpointed. The checkpoint is returned to
13 the client in a Task message.
- 14 GetTaskStatus:
15 Requests the status of a particular task. The server will send a TaskStatus
16 message in response.
- 17 RemoveTask:
18 Requests the removal of the named task.
- 19 ExecutableElement:
20 Register an executable element with the server. The actual executable element
21 data must be downloaded over the data channel.
- 22 DataElement:
23 Register a data element with the server. The actual data element data must be
24 downloaded over the data channel.
- 25 CloseSession:
26 Close current session.
- 27 Error:
28 Report an error condition.
- 29 ExternalDelivery:

- 1 Alert the recipient that a message is waiting at a specified URL.
- 2 Heartbeat:
- 3 Alert the recipient that the sending entity is still active.
- 4
- 5 Server-to-Provider messages
- 6 Task:
- 7 Assigns a task to a provider.
- 8 GetTaskStatus:
- 9 Requests a task status for a particular task. Provider responds with a
- 10 TaskStatus message.
- 11 GetTask:
- 12 Requests a checkpoint for a particular task. The provider responds with a Task
- 13 message containing the checkpoint.
- 14 Config:
- 15 Sent to change the providers server-controlled configuration settings.
- 16 GetContents:
- 17 Request contents of providers message queue.
- 18 WorkSummary:
- 19 Send current node statistics to CE for display to the user.
- 20 GetCacheContents:
- 21 Request contents of provider's cache. Provider responds with a CacheContents
- 22 message.
- 23 DisplayContent:
- 24 Send new display content to CE.
- 25 Upgrade:
- 26 Send a software upgrade to CE.
- 27 Error:
- 28 Report an error condition.
- 29 ExternalDelivery:

- 1 Alert the recipient that a message is waiting at a specified URL.
- 2
- 3 Provider-to-Server messages
- 4 NodeStatus:
- 5 Reports a change in the status of the provider node.
- 6 NodeProfile:
- 7 Reports profile information for node (e.g., machine type, available disk,
8 available memory, connection speed, operating system type).
- 9 GetConfig:
- 10 Request that the server send a Config message to set the providers server-
11 controlled configuration settings.
- 12 CacheContents:
- 13 Reports the contents of the providers element cache. This information is useful
14 in task scheduling.
- 15 Contents:
- 16 Reports the contents of the task queue.
- 17 Task Status:
- 18 Reports the status of a task. There can be both intermediate and final task
19 statuses. Intermediate task statuses may contain partial results. Final task
20 statuses contain the results of the task. Task statuses are also used to report
21 errors in the execution of a task.
- 22 Task:
- 23 The provider sends the task message out in response to a GetTask message
24 from the provider. In this case, the Task message is a checkpoint of the state of
the identified task.
- 25 Error:
- 26 Report an error condition.
- 27 ExternalDelivery:
- 28 Alert the recipient that a message is waiting at a specified URL.
- 29 Heartbeat:

1 Alert the recipient that the sending entity is still active.

2

3 **Server-to-Client messages**

4 **Task:**

5 The server sends the task message out in response to a GetTask message from
6 the client. In this case, the Task message is a checkpoint of the state of the
7 identified task.

8 **Task Status:**

9 Reports the status of a task. This could be generated in response to a task status
10 from the provider, in response to a GetStatus from the client, or due to an
11 internal server condition.

12 **Config:**

13 Sent to change the clients server-controlled configuration settings.

14 **JobProperties:**

15 Report current job properties.

16 **Contents:**

17 Return contents of job or client global area.

18 **ExecutableElement:**

19 Return information on a specific executable element. Sent in response to a
20 GetExecutableElement message.

21 **DataElement:**

22 Return information on a specific data element. Sent in response to a
23 GetDataElement message.

24 **Error:**

25 Report an error condition.

26 **ExternalDelivery:**

27 Alert the recipient that a message is waiting at a specified URL.

28

1 The present invention can also be employed in other manners, such as a method of
2 marketing computers by offering incentives to computer customers that agree to operate a
3 compute engine (CE) on the computers and having the CE utilize idle computational
4 processing power on the computers. Incentives can include, but are not limited to free
5 computer use, free ISP service, discounted computer sales price, discount computer lease
6 price, a sales rebate, periodic rebates, and usage fees. The CE can also be utilized to deliver
7 “pushed” content, such as advertising, to these computer customers via a display window of
8 said computer’s graphic interface or via said computer’s sound output.

9 A system and method for the establishment and utilization of networked idle
10 computational processing power has been illustrated. It will be appreciated by those skilled in
11 the art that other variations in, for example, the calculation of power, and the methods for
12 compensating providers can be accomplished without departing from the scope of the
13 invention as disclosed.

14

1 We Claim:

- 2 1. A system for using computer excess computational power comprising:
3 at least one client computer adapted to create programs for execution;
4 a network connected to the client computer;
5 a central task server (CTS) connected to the network, adapted to receive the programs
6 from the client;
7 a plurality of provider computers for providing excess computational power connected
8 to the CTS via network and adapted to receive the programs from the CTS, the CTS thereby
9 providing computational power upon demand to the at least one client computer.
- 10 2. The system for using excess computational power of claim 1 wherein the at least one
11 client computer further comprises:
12 a client application programming interface (API) for allowing the client computer to
13 specify the compute capacity needed in the form of a virtual cluster in which the programs are
14 to be run.
- 15 3. The system for using excess computational power of claim 2 wherein the parameters
16 to be specified are at least computer power, disk space required, and reliability of the
17 computer power.
- 18 4. The system for using excess computational power of claim 2 wherein the at least one
19 client computer further comprises a run time interface, the run time interface further
20 comprising a set of function calls used by the client computer for getting parameters from the
21 run time environment, for getting data elements, for querying the CTS for task status, and for
22 performing checkpointing at various places in the program.
- 23 5. The system for using excess computational power of claim 4 wherein the run time
24 interface further comprises instructions for creating jobs.
- 25 6. The system for using excess computational power of claim 5 wherein the run time
26 interface further comprises instructions for creating tasks within jobs.
- 27 7. The system for using excess computational power of claim 5 wherein the run time
28 interface further comprises instructions for querying against the status of jobs.

- 1 8. The system for using excess computational power of claim 5 wherein the run time
2 interface further comprises instructions for attaching the client computer of a job in progress
3 and for checking on the status of a job while the job is in progress.
- 4 9. The system for using excess computational power of claim 1 wherein the CTS further
5 comprises instructions for charging the at least one client for computational power used
6 during the execution of the programs.
- 7 10. The system for using excess computational power of claim 9 wherein the charging of
8 the client is made when the task is launched.
- 9 11. The system for using excess computational power of claim 10 wherein a task is
10 launched when it is scheduled by the CTS to be run by at least one provider computer.
- 11 12. The system for using excess computational power of claim 1 wherein the CTS further
12 comprises instructions for monitoring a virtual node specified by the client for failure.
- 13 13. The system for using excess computational power of claim 9 wherein charging of
14 clients is based upon the sum of the power of the provider computers running the client tasks.
- 15 14. The system for using excess computational power of claim 9 wherein the CTS further
16 comprises instructions for making payments to providers based upon the available
17 computational power the provider's computer.
- 18 15. The system for using excess computational power of claim 14 wherein the available
19 computational power is a function of CPU availability, CPU speed and bandwidth
20 availability.
- 21 16. The system for using excess computational power of claim 15 wherein CPU speed is
22 measured by running benchmark tests of the CPU periodically.
- 23 17. The system for using excess computational power of claim 1 wherein each of the
24 plurality of provider computer further comprises a general purpose compute engine (CE) for
25 executing programs sent from the central server.
- 26 18. The system for using excess computational power of claim 17 wherein the CE further
27 comprises instructions for allowing an owner of the provider computer to set the availability
28 parameters of the provider computer.

19. The system for using excess computational power of claim 18 wherein the availability parameters comprise the amount of RAM available for the programs, the amount of disk space available for the programs and the times during the day, if any, that the provider computer is available to execute the programs.

20. The system for using excess computational power of claim 18 wherein the CE further comprises instructions for opening a window on the provider computer when the programs are being run.

21. The system for using excess computational power of claim 18 wherein the CE further comprises instructions for displaying content from the central server in the window.

22. The system for using excess computational power of claim 18 wherein the CE further comprises instructions for executing the programs only when excess computational power is detected on the provider computer.

23. The system for using excess computational power of claim 9 wherein the amount charged to the client I based on the following formula:

$$\text{Price per hour for computation} = 4.5 + 0.296P^{1.155},$$

where P is the sum of all of the power in active nodes being used by the CTS to fulfill the computational requirements specified by the client.

24. The system for using excess computational power of claim 23 wherein the client is also charged for data transfer during the course of executing the particular tasks of the client based upon the following formula:

$$\text{Total Cost of Data Transfer} = g(d_i + d_o) \left(\frac{\ln(1-r)}{\ln(1-\pi)} + 1 \right),$$

where d_i is the quantity (in GB) of data transferred from the client to the server, d_o is the quantity of data transferred from the server to the client, where g is the transfer cost per communication path among the client computer, the provider computer and the CTS per GB, r is a requested node reliability, and π is an average provider computer reliability.

25. The system for using excess computational power of claim 24 wherein the client is also charged an hourly premium for requesting reliability greater than a set baseline reliability given by the formula:

$$\text{Price per hour for premium reliability} = kCN \left(\frac{\ln(1-r)}{\ln(1-R)} - 1 \right)$$

where k is a baseline price per MHz-hour, C is the CPU speed requested of a virtual node requested by the client, N is the number of virtual nodes running, r is the requested reliability, and R is the baseline reliability.

26. The system for using excess computational power of claim 1 wherein the owners of the provider computers are paid a fixed amount per unit work for providing excess computational power according to the following formula:

$$\frac{T \ln(1-\pi)}{\ln(1-r)} \text{ per GHz-hour, where}$$

T is a rate, π is an average availability, and r is a target reliability.

27. A method for using computer excess computational power comprising:
 adapting at least one client computer to create programs for execution;
 connecting the client computer to a network;
 connecting a central task server (CTS) to the network, wherein said CTS is adapted to receive the programs from the client;
 connecting a plurality of provider computers for providing excess computational power to the network and adapting said provider computers to receive the programs from the central server, the plurality of provider computers providing computational power upon demand.

28. The method for using excess computational power of claim 27 further comprising supplying the at least one client computer with a client application programming interface (API) for allowing the client computer to specify the parameters under which the programs are to be run.

1 29. The method for using excess computational power of claim 28 wherein the
2 parameters to be specified are at least computer power, disk space required, and reliability of
3 the computer power.

4 30. The method for using excess computational power of claim 28, further comprising
5 providing the at least one client computer with a run time interface, the run time interface
6 including a set of function calls used by the client computer for getting parameters from the
7 run time environment, for getting data elements, for querying the CTS for task status, and for
8 performing checkpointing at various places in the program.

9 31. The method for using excess computational power of claim 30, further comprising
10 providing the run time interface with instructions for creating jobs.

11 32. The method for using excess computational power of claim 31, further comprising
12 providing the run time interface with instructions for creating tasks within jobs.

13 33. The method for using excess computational power of claim 31, further comprising
14 providing the run time interface with instructions for querying against the status of jobs.

15 34. The method for using excess computational power of claim 31, further comprising
16 providing the run time interface with instructions for attaching the client computer of a job in
17 progress and for checking on the status of a job while the job is in progress.

18 35. The method for using excess computational power of claim 27, further comprising
19 providing the CTS with instructions for charging the at least one client for excess
20 computational power used during the execution of the programs.

21 36. The method for using excess computational power of claim 35 wherein the charging of
22 the client is made when the task is launched.

23 37. The method for using excess computational power of claim 36 including launching a
24 task when it is scheduled by the CTS to be run by at least one provider computer.

25 38. The method for using excess computational power of claim 27, further comprising
26 providing the CTS with instructions for monitoring a virtual node specified by the client for
27 failure.

28 39. The method for using excess computational power of claim 35 wherein charging of
29 clients is based upon the sum of the power of the provider computers running the client tasks.

1 40. The method for using excess computational power of claim 35, further comprising
2 providing the CTS with instructions for making payments to providers based upon the
3 available computational power the provider's computer.

4 41. The method for using excess computational power of claim 40 wherein the available
5 computational power is determined as a function of CPU availability, CPU speed and
6 bandwidth availability.

7 42. The method for using excess computational power of claim 41 wherein CPU speed is
8 measured by running benchmark tests of the CPU periodically.

9 43. The method for using excess computational power of claim 27, further comprising
10 providing each of the plurality of provider computer with a general purpose compute engine
11 (CE) for executing programs sent from the central server.

12 44. The method for using excess computational power of claim 43, further comprising
13 providing the CE with instructions for allowing an owner of the provider computer to set the
14 availability parameters of the provider computer.

15 45. The method for using excess computational power of claim 44 wherein the availability
16 parameters comprise the amount of RAM available for the programs, the amount of disk
17 space available for the programs and the times during the day, if any, that the provider
18 computer is available to execute the programs.

19 46. The method for using excess computational power of claim 44, further comprising
20 providing the CE with instructions for opening a widow on the provider computer when the
21 programs are being run.

22 47. The method for using excess computational power of claim 44, further comprising
23 providing the CE with instructions for displaying content from the central server in the
24 window.

25 48. The method for using excess computational power of claim 44, further comprising
26 providing the CE with instructions for executing the programs only when excess
27 computational power is detected on the provider computer.

49. The method for using excess computational power of claim 35 wherein the amount charged to the client I is computed based on the following formula:

$$\text{Price per hour for computation} = 4.5 + 0.296P^{1.155},$$

where P is the power of the computer in Mhz.

50. The method for using excess computational power of claim 49, further comprising charging the client for data transfer during the course of executing the particular tasks of the client based upon the following formula:

$$\text{Total Cost of Data Transfer} = g(d_i + d_o) \left(\frac{\ln(1-r)}{\ln(1-\pi)} + 1 \right),$$

where d_i is the quantity (in GB) of data transferred from the client to the server, d_o is the quantity of data transferred from the server to the client, where g is the transfer cost per communication path among the client computer, the provider computer and the CTS per GB, r is a requested node reliability, and π is an average provider computer reliability.

51. The method for using excess computational power of claim 50, further comprising charging the client an hourly premium for requesting reliability greater than a set baseline reliability given by the formula:

$$\text{Price per hour for premium reliability} = kCN \left(\frac{\ln(1-r)}{\ln(1-R)} - 1 \right)$$

where k is a baseline price per MHz-hour, C is the CPU speed requested of a virtual node requested by the client, N is the number of virtual nodes running, r is the requested reliability, and R is the baseline reliability.

52. The method for using excess computational power of claim 27, further comprising paying the owners of the provider computers a fixed amount per unit work for providing excess computational power according to the following formula:

$$\frac{T \ln(1-\pi)}{\ln(1-r)} \text{ per GHz-hour, where}$$

1 T is a rate, π is an average availability, and r is a target reliability.

2 53. A method for the establishment and utilization of networked idle computational
3 processing power comprising:

4 loading a client API on to client computers to allow client jobs to be broken into
5 independent tasks.

6 loading compute engine (CE) software onto provider computers to execute tasks using
7 idle computational processing power of said provider computers and using a centralized task
8 server (CTS) that is at least intermittently networked to said client computers at any time
9 selected by said clients and intermittently connected to provider computers to exchange tasks
10 and results.

11 54. The method of claim 53, wherein said CE runs unobtrusively and securely on the
12 provider computer.

13 55. The method of claim 53, wherein a task scheduler in said CTS matches
14 computational power requested by clients to provider resources.

15 56. The method of claim 53, wherein CE operates the provider computer to
16 periodically contact the CTS to retrieve tasks and control messages and/or send results.

17 57. The method of claim 53, wherein the client API periodically contacts the CTS to
18 retrieve results.

19 58. The method of claim 53, wherein the CE displays a graphic window when running
20 tasks to display processing and status information.

21 59. The method of claim 58, further comprising display of advertising on other
22 targeted messages by the CE graphic window.

23 60. The method of claim 58, wherein the window can be made to disappear based on
24 configurable mouse coordinates.

25 61. A system for the established and utilization of networked idle computational
26 processing power comprising:

27 at least one client computer with a client API that allows client jobs to be broken into
28 independent tasks;

1 at least one provider computer operating compute engine (CE) software that executes
2 tasks to produce results when said provider computer processor is idle

3 at least one centralized task server (CTS) at least intermittently connected to said
4 client at any time of said client's choosing and at least intermittently connected to provider
5 computers to exchange tasks and results.

6 62. The system of claim 61, wherein said CTS comprises redundant, fault tolerant and
7 scalable architecture.

8 63. A method of marketing computers, comprising:
9 offering incentives to computer customers that agree to operate a compute engine
10 (CE) on said computers; and
11 having said CE utilize idle computational processing power on said computers to
12 perform tasks requested by an intermediary that markets said idle computational
13 processing power.

14 64. The method of marketing computers of claim 63, wherein said incentives are
15 selected from the group consisting of free computer use, free ISP service,
16 discounted computer sales price, discount computer lease price, a sales rebate,
17 periodic rebates, and usage fees.

18 65. The method of marketing computers of claim 63, wherein said CE is also utilized
19 to deliver content to said computer customers via a display window of said
20 computer's graphic interface or via said computer's sound output.

21 66. The method of marketing computers of claim 65, wherein said content is
22 advertising.

23 67. A method of creating and intermediating a market for idle computational
24 processing power comprising:
25 aggregating computational tasks from clients;
26 distributing tasks to provider computers for execution using idle computational
27 processing power of said provider computers;
28 collecting results from said provider computers and delivering said results to said
29 clients;

1 charging said clients a fee; and

2 compensating said providers.

3 68. The method of creating and intermediating a market for idle computational
4 processing power of claim 67 wherein API software on client computers provides
5 tasks and compute engine (CE) software on said provider computers executes said
6 tasks and provides said results.

7 69. The system of claim 61, further comprising a communication protocol for
8 governing communication between the API, CTS, and CE, said protocol having
9 separate message and data channels; wherein the message channel is a bi-
10 directional channel between clients and the server that is used to pass high-level
11 messages that facilitate distributed processing and the data channel is a channel
12 used to pass large blocks of data elements and executable elements.

13 70. The system of claim 61, wherein the client API supports jobs of arbitrary size by
14 being programmed to monitor only task results.

15 71. The system of claim 61, further comprising a runtime API which defines how
16 tasks interact with their environment.

17 72. The system of claim 61, wherein the client API contains an interface for task
18 checkpointing.

19 73. The system of claim 61, wherein the CE supports the queuing of tasks, task
20 prioritization, and preemption of tasks based on priority.

21 74. The system of claim 61, wherein the CTS and CE together record computer CPU
22 and communications line signatures.

23 75. The system of claim 74, wherein the signatures are historical usage patterns for use
24 in task scheduling.

25 76. The system of claim 61, further comprising system-wide secure communications,
26 client anonymity, and obfuscation of client code.

27 77. The system of claim 61, further comprising means for pricing jobs and scheduling
28 tasks based on delivering clients a virtual cluster comprised of virtual nodes, said

1 virtual nodes implemented by launching the same task to multiple, complementary
2 nodes that together provide a given expected level of power and reliability.

3 78. The system of claim 69, wherein the CE

4 79. A general-purpose compute engine (CE) for execution of distributed tasks using
5 the excess computational capacity of provider machines comprising:

6 compute engine software having a virtual machine for executing tasks; and

7 means for orchestrating the execution of tasks via said virtual machine in a manner

8 that keeps said provider machine substantially insulated from client tasks.

9 80. The general-purpose compute engine of claim 79, wherein the CE is configured to
10 run from the group consisting of constantly and only when said provider computer
11 is idle.

12 81. The system for using excess computational power of claim 13 wherein the client is
13 charged and enhanced rate for power used above a set threshold.

14

1/8




Figure 1

2 / 8

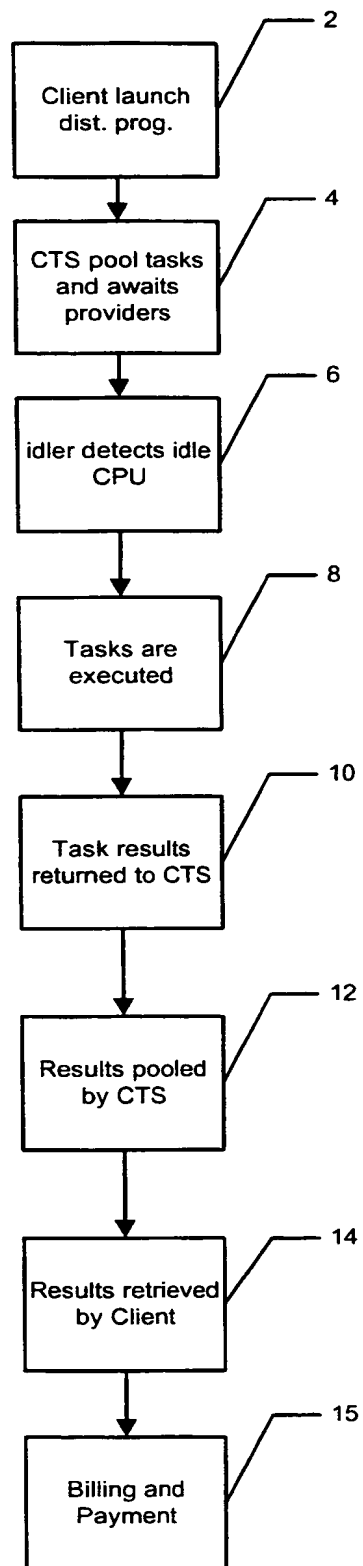



Figure 2

Figure 3



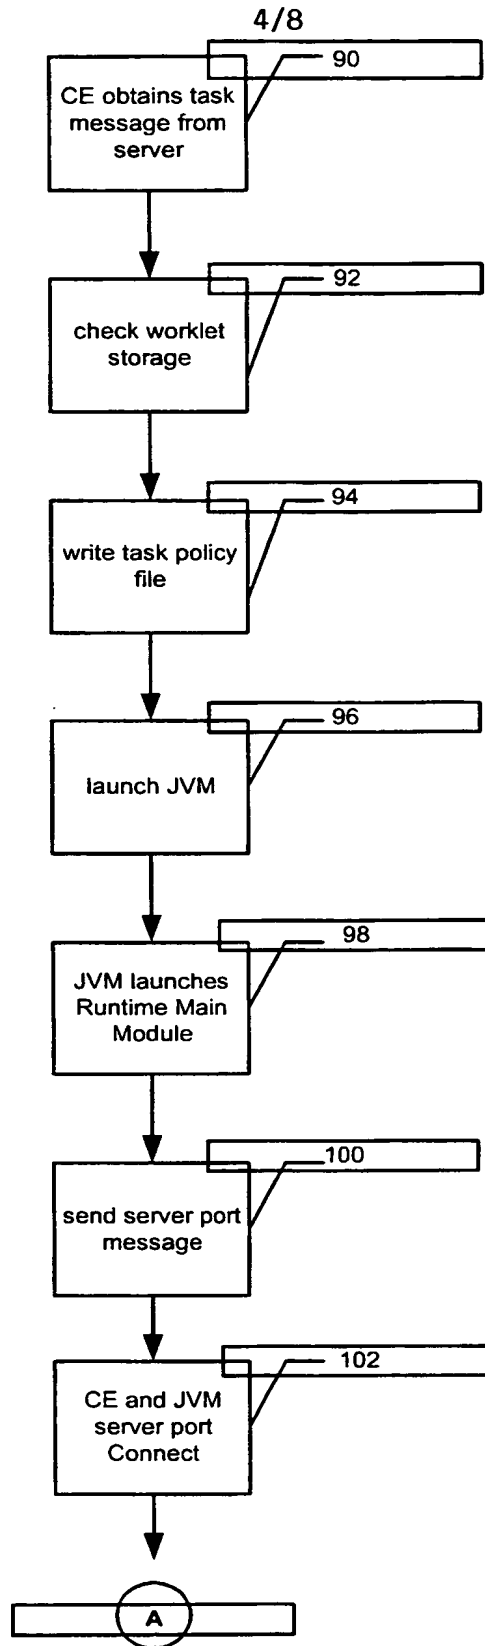


Figure 4A

5/8

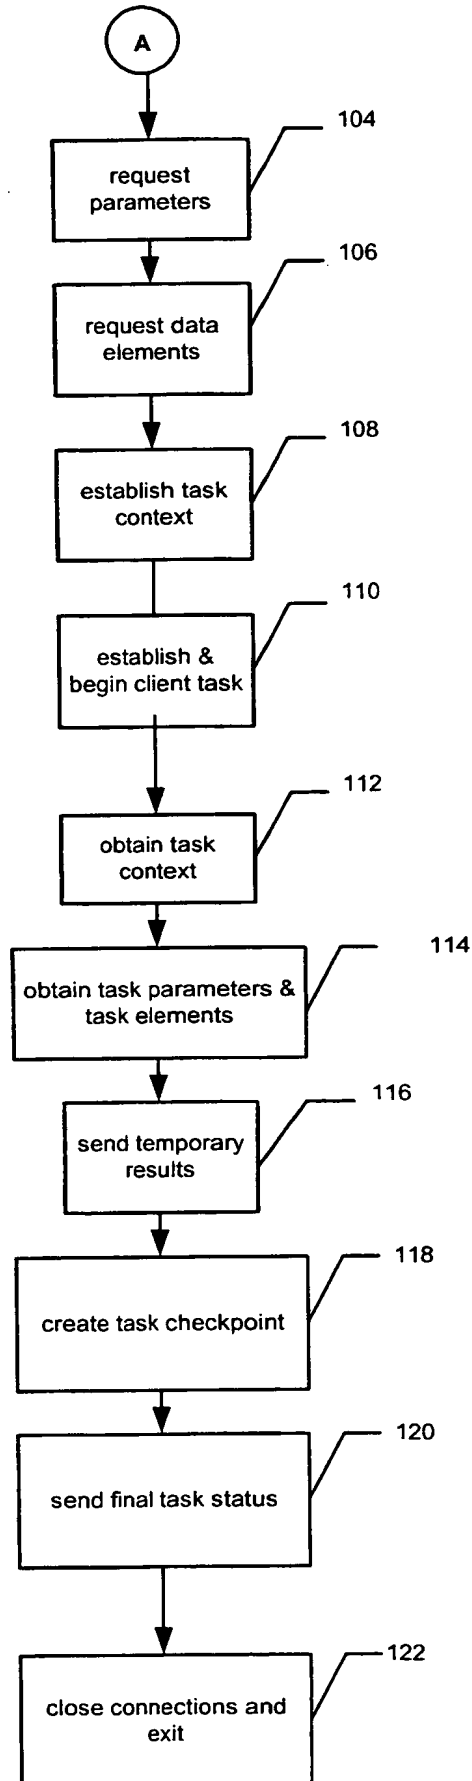


Figure 4B

6 / 8

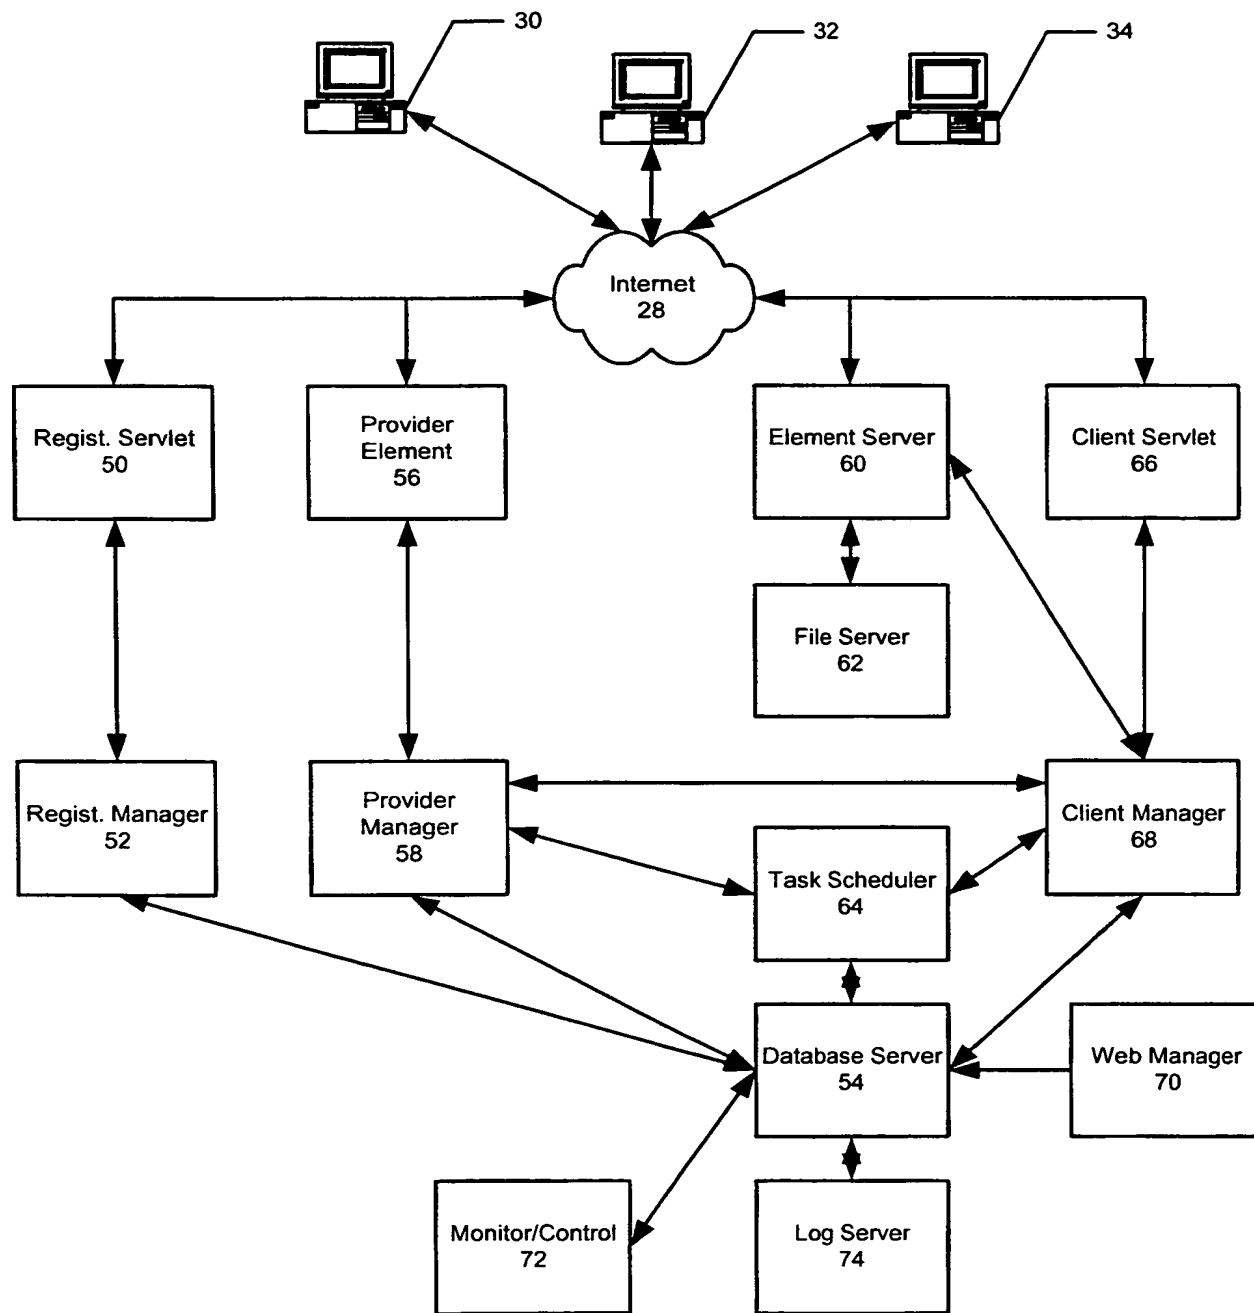



Figure 5

7/8

Figure 6



8 / 8




Figure 7




Figure 8

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
1 March 2001 (01.03.2001)

PCT

(10) International Publication Number
WO 01/14961 A3

(51) International Patent Classification⁷: **G06F 9/50**

(21) International Application Number: **PCT/US00/24336**

(22) International Filing Date: **28 August 2000 (28.08.2000)**

(25) Filing Language: **English**

(26) Publication Language: **English**

(30) Priority Data:
60/150,766 26 August 1999 (26.08.1999) US
60/210,334 12 June 2000 (12.06.2000) US

(71) Applicant: **PARABON COMPUTATION [US/US];**
Suite 1000, 3930 Walnut Street, Fairfax, VA 22030 (US).

(72) Inventors: **ARMENTROUT, Steven, L.;** 11879 St. Trinians Ct., Reston, VA 20191 (US). **O'CONNOR, James;** 11625 Ayreshire Road, Oakton, VA 22124 (US). **GANNON, James;** 3817 Inverness Road, Fairfax, VA 22030 (US). **SLETTEN, Brian;** Apt. 906, 900 N. Stuart Street, Arlington, VA 22203 (US). **CIER, Sean;** Apt. 12, 11800 Federalist Way, Fairfax, VA 22030 (US). **CARLSON, Sarah;** 13236 Memory Lane, Fairfax, VA 22033 (US). **DAVIS, Jonathan;** Apt. 312, 202 Columbia Road NW, Washington, DC 20009 (US). **DUPERTUIS, Greg;** 7888 Leroux Lane, Manassas, VA 20112 (US). **MCCLOUGHLIN, Scott;** 2827 28th Street, #14, Washington, DC 20001 (US). **DAVIES, Antony;** 11733 Stuart Mill Road, Oakton, VA 22124 (US).

(74) Agents: **ROBERTS, Jon, L. et al.;** Roberts Abokhair & Mardula, LLC, Suite 1000, 11800 Sunrise Valley Drive, Reston, VA 20191 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).


Published:

— with international search report

(88) Date of publication of the international search report:
2 August 2001

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: **SYSTEM AND METHOD FOR THE ESTABLISHMENT AND UTILIZATION OF NETWORKED IDLE COMPUTATIONAL PROCESSING POWER**



(57) Abstract: A distributed computing platform using the idle computational processing power of a plurality of provider computers is disclosed. At least one networked server collects tasks from client computers, schedules and distributes the tasks to networked provider computers, and collects and returns results to client computers. A client API forms tasks and collects results. A compute engine operates on the provider computers to communicate with the server and execute tasks using idle computational power.

WO 01/14961 A3

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 00/24336

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F9/50

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, INSPEC, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	LONDON, SHMULIK: "POPCORN - A Paradigm for Global-Computing" June 1998 (1998-06), THE HEBREW UNIVERSITY OF JERUSALEM (INSTITUTE OF COMPUTER SCIENCE), JERUSALEM, ISRAEL XP002159919 the whole document --- -/--	1-81

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *Z* document member of the same patent family

Date of the actual completion of the international search

9 February 2001

Date of mailing of the international search report

22/02/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Bijn, K

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 00/24336

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	CHAO-JU HOU ET AL: "IMPLEMENTATION OF DECENTRALIZED LOAD SHARING IN NETWORKED WORKSTATIONS USING THE CONDOR PACKAGE" JOURNAL OF PARALLEL AND DISTRIBUTED COMPUTING,US,ACADEMIC PRESS, DULUTH, MN, vol. 40, no. 2, 1 February 1997 (1997-02-01), pages 173-184, XP000682281 ISSN: 0743-7315 page 174, left-hand column, line 35 -page 175, right-hand column, last line ----	1-3,17, 27-29, 43,63, 79,80
A	WALDSPURGER C A ET AL: "Spawn: a distributed computational economy" IEEE TRANSACTIONS ON SOFTWARE ENGINEERING,US,IEEE INC. NEW YORK, vol. 18, no. 2, February 1992 (1992-02), pages 103-117, XP002124500 ISSN: 0098-5589 the whole document ----	1-81
A	MUTKA M W: "ESTIMATING CAPACITY FOR SHARING IN A PRIVATELY OWNED WORKSTATION ENVIRONMENT" IEEE TRANSACTIONS ON SOFTWARE ENGINEERING,US,IEEE INC. NEW YORK, vol. 18, no. 4, 1 April 1992 (1992-04-01), pages 319-328, XP000328663 ISSN: 0098-5589 the whole document ----	67,74
A	TANDIARY F ET AL: "Batrung: utilizing idle workstations for large scale computing" IEEE PARALLEL AND DISTRIBUTED TECHNOLOGY: SYSTEMS AND APPLICATIONS,XX,XX, vol. 4, no. 2, 21 June 1996 (1996-06-21), pages 41-48, XP002124512 ISSN: 1063-6552 the whole document -----	1,17,27, 43,53, 54,56,61

THIS PAGE BLANK (USPTO)